

CLAIMS:

Having thus described our invention, what we claim as new, and desire to secure by Letters Patent is:

1 / 1. In a group-based multicast messaging system implementing
2 dedicated logger member for logging all multicast messages sent
3 in said system, a method for reliably delivering messages from
4 senders to receivers of said group, said method comprising steps
5 of:

6 a) a receiver detecting one or more missing messages from a
7 sequence of multicast messages sent to members of said group;

8 b) soliciting retransmissions of missing messages to another
9 member or logger of said group;

10 c) determining said receiver's missing messages as fresh or
11 stale, and one of: enabling repair of fresh missing messages by
12 said another member in a first message recovery phase or,
13 enabling repair of stale missing messages by a logger in a second
14 message recovery phase, wherein reliable delivery of messages in
15 said multicast messaging system is ensured.

1 2. The method as claimed in Claim 1, wherein each multicast
2 message is assigned a sequence number, said step a) of detecting
3 a missing message includes determining a gap in sequence numbers
4 of messages received by a member.

1 3. The method as claimed in Claim 2, wherein said missing
2 message status is determined by comparing a time difference
3 between a current time and the time of actual sending of said
4 missing message in said system against a fixed time threshold,

5 wherein when said time threshold has not been exceeded,
6 assigning said fresh status to said missing message, otherwise,
7 assigning said stale status to said missing message.

1 4. The method as claimed in Claim 2, wherein said step b) of
2 soliciting retransmission includes:

3 generating a gossip message comprising a retransmission
4 request of said missing message;

5 randomly selecting a member of said group; and,

6 sending said gossip message to said randomly selected
7 member.

1 5. The method as claimed in Claim 4, wherein said gossip message
2 comprises a negative gossip comprising a member's missing
3 messages.

4 6. The method as claimed in Claim 4, further comprising the
5 steps of:

6 collecting latest state information for all senders in said
7 group in a gossip message, said state information including
8 latest sequence numbers transmitted by said senders; and,

9 propagating said state information in said gossip message.

1 7. The method as claimed in Claim 6, further including the step
2 of applying a selective hash function to said state information
3 to generate a hash signature, and including said hash signature
4 in said gossip message.

1 8. The method as claimed in Claim 7, wherein said step b) of
2 soliciting retransmission includes:

3 generating a retransmission request for said missing
4 message;
5 randomly selecting a logger of said group; and,
6 sending said retransmission request to said randomly
7 selected logger.

1 9. The method as claimed in Claim 8, wherein each logger member
2 generates acknowledgment message for receipt by a sender member
3 in response to logging a message received from said sender, said
4 method further comprising the step of enabling a sender to
5 further derive an acknowledgment from a received gossip message.

1 10. The method as claimed in Claim 9, wherein each member
2 maintains a hash signature of its own state, said deriving step
3 including the step of: comparing said hash signature of a
4 member's state with a hash signature embedded in said gossip
5 message to determine which messages a gossipier has received for
6 treatment as said acknowledgment.

1 11. The method as claimed in Claim 8, further including the step
2 of maintaining membership status of group members associated with
3 multicast messages, said state information propagated in said
4 gossip message further including embedded group membership status
5 for said group members.

1 12. The method as claimed in Claim 11, wherein each logger
2 performs steps of:
3 ascertaining views of group memberships for each message
4 from said state information;
5 determining if all members of said group has received the
6 logged message; and,

7 initiating garbage collection of said logged messages
8 determined as being received by all said members.

1 13. The method as claimed in Claim 12, wherein a logger member
2 maintains a pending acknowledgment list comprising current group
3 members associated with each new message received, said step of
4 determining if all members of said group has received the logged
5 message includes:

6 deriving acknowledgments from state information included in
7 said gossip messages; and,
8 verifying receipt of the logged message for all members in
9 said pending acknowledgment list.

10 14. A reliable recovery system for a group-based multicast
11 messaging system implementing dedicated logger member for logging
12 all multicast messages sent by system group members, said
13 recovery system comprising:

14 mechanism in system group members for detecting one or more
15 missing messages from a sequence of received multicast messages
16 sent via said system and, for communicating requests for messages
17 determined missing by a group member to another group member or
logger;

18 mechanism for determining a status of a missing message
19 indicated in a request as one of a fresh missing message and a
20 stale missing message;

21 first phase repair mechanism for enabling repair of missing
22 fresh messages by another group member; and,

23 second phase repair mechanism for enabling repair of missing
24 stale messages by a logger member, wherein reliable delivery of
25 messages in said multicast messaging system is ensured.

1 15. The reliable recovery system as claimed in Claim 14, wherein
2 a group member assigns a sequence number to each message, said
3 detecting mechanism including a mechanism for determining a gap
4 in sequence numbers of messages received by a member.

1 16. The reliable recovery system as claimed in Claim 15, wherein
2 said missing message status is determined by comparing a time
3 difference between a current time and the time of actual sending
4 of said missing message in said system against a fixed time
5 threshold,

6 wherein when said time threshold has not been exceeded, said
7 message is assigned said fresh status, and otherwise, is assigned
8 said stale status.

1 17. The reliable recovery system as claimed in Claim 15, wherein
2 said mechanism for communicating requests comprises:

3 mechanism for generating a gossip message comprising said
4 detected one or more missing messages;

5 mechanism for randomly selecting a member of said group;
6 and,

7 transmit device for sending said gossip message to said
8 randomly selected member.

1 18. The reliable recovery system as claimed in Claim 17, wherein
2 said gossip message comprises a negative gossip comprising a
3 member's missing messages.

1 19. The reliable recovery system as claimed in Claim 17, wherein
2 said mechanism for generating a gossip message further comprises:
3 a mechanism for collecting latest state information for all
4 senders in said group and embedding said state information in a

5 gossip message, said state information including latest sequence
6 numbers transmitted by said senders.

1 20. The reliable recovery system as claimed in Claim 19, further
2 including mechanism for applying selective hash function to said
3 state information to generate a hash signature for embedding in
4 said gossip message.

1 21. The reliable recovery system as claimed in Claim 20, wherein
2 said mechanism for communicating requests comprises:

3 mechanism for generating a retransmission request for said
4 missing message;

5 mechanism for randomly selecting a logger of said group;
6 and,

7 transmit device for sending said retransmission request to
8 said randomly selected logger.

9 22. The reliable recovery system as claimed in Claim 21, wherein
10 each logger member includes mechanism responsive to a received
11 message for generating an acknowledgment message for receipt by a
12 sender member after logging said message, said system further
13 comprising mechanism for enabling sender derivation of an
14 acknowledgment from a received gossip message.
15
16

1 23. The reliable recovery system as claimed in Claim 22, wherein
2 each member maintains a hash signature of its own state, said
3 deriving mechanism including mechanism for comparing said hash
4 signature of a member's state with a hash signature embedded in
5 said gossip message to determine which messages a gossipper has
6 received for treatment as said acknowledgment.

1 24. The reliable recovery system as claimed in Claim 20, further
2 including implementation of a relaxed group membership protocol
3 for maintaining membership status of group members associated
4 with multicast messages, said state information further including
5 embedded group membership status of said group members.

1 25. The reliable recovery system as claimed in Claim 24, wherein
2 each logger includes mechanism for ascertaining views of group
3 memberships for each message from received state information and
4 determining if all members of said group has received the logged
5 message.

1 26. The reliable recovery system as claimed in Claim 25, wherein
2 each logger further includes garbage collection mechanism for
3 removing said logged messages determined as being received by all
4 said members.

1 27. The reliable recovery system as claimed in Claim 26, wherein
2 a logger member maintains a pending acknowledgment list
3 comprising current group members associated with each new message
4 received, said mechanism for determining if all members of said
5 group has received the logged message including mechanism for
6 deriving acknowledgments from state information included in said
7 gossip messages and verifying receipt of the logged message for
8 all members in said pending acknowledgment list.

1 28. A program storage device readable by a machine, tangibly
2 embodying a program of instructions executable by the machine to
3 perform method steps for reliably delivering messages from
4 senders to receivers of a group-based multicast messaging system,
5 the system implementing a dedicated logger member for logging all

6 multicast messages sent in said system, said method steps
7 comprising:

8 a) a receiver detecting one or more missing messages from a
9 sequence of multicast messages sent to members of said group;

10 b) soliciting retransmissions of missing messages to another
11 member or logger of said group;

12 c) determining said receiver's missing messages as fresh or
13 stale, and one of: enabling repair of fresh missing messages by
14 said another member in a first message recovery phase or,
15 enabling repair of stale missing messages by a logger in a second
16 message recovery phase, wherein reliable delivery of messages in
17 said multicast messaging system is ensured.

29. The program storage device readable by a machine according
to Claim 28, wherein each multicast message is assigned a
sequence number, said step a) of detecting a missing message
includes determining a gap in sequence numbers of messages
received by a member.

30. The program storage device readable by a machine as claimed
in Claim 29, wherein said missing message status is determined by
comparing a time difference between a current time and the time
of actual sending of said missing message in said system against
a fixed time threshold, wherein when said time threshold has not
been exceeded, assigning said fresh status to said missing
message, otherwise, assigning said stale status to said missing
message.

31. The program storage device readable by a machine as claimed
in Claim 29, wherein said step b) of soliciting retransmission
includes:

4 generating a gossip message comprising a retransmission
5 request of said missing message;
6 randomly selecting a member of said group; and,
7 sending said gossip message to said selected member.

1 32. The program storage device readable by a machine as claimed
2 in Claim 31, wherein said gossip message comprises a negative
3 gossip comprising a members missing messages.

1 33. The program storage device readable by a machine as claimed
2 in Claim 31, further comprising the steps of collecting latest
3 state information for all senders in said group in a gossip
4 message, said state information including latest sequence numbers
5 transmitted by said senders, and propagating said state
6 information in said gossip message.

1 34. The program storage device readable by a machine as claimed
2 in Claim 33, further including the step of applying a selective
3 hash function to said state information to generate a hash
4 signature, and including said hash signature in said gossip
5 message.

1 35. The program storage device readable by a machine as claimed
2 in Claim 34, wherein said step b) of soliciting retransmission
3 includes:

4 generating a retransmission request for said missing
5 message;
6 randomly selecting a logger of said group; and,
7 sending said retransmission request to said randomly
8 selected logger.

1 36. The program storage device readable by a machine as claimed
2 in Claim 35, wherein each logger member generates acknowledgment
3 message for receipt by a sender in response to logging a message
4 received from said sender, said method steps further comprising
5 the step of enabling a sender to further derive an acknowledgment
6 from a received gossip message.

1 37. The program storage device readable by a machine as claimed
2 in Claim 36, wherein each member maintains a hash signature of
3 its own state, said deriving step including the step of:
4 comparing said hash signature of a member's state with a hash
5 signature embedded in said gossip message to determine which
6 messages a gossipier has received for treatment as said
acknowledgment.

1 38. The program storage device readable by a machine as claimed
2 in Claim 35, further including the step of maintaining membership
3 status of group members associated with multicast messages, said
4 state information propagated in said gossip message further
5 including embedded group membership status for said group
6 members.

1 39. The program storage device readable by a machine as claimed
2 in Claim 38, wherein each logger performs steps of:
3 ascertaining views of group memberships for each message
4 from said state information;
5 determining if all members of said group has received the
6 logged message; and,
7 initiating garbage collection of said logged messages
8 determined as being received by all said members.

1 40. The program storage device readable by a machine as claimed
2 in Claim 39, wherein a logger member maintains a pending
3 acknowledgment list comprising current group members associated
4 with each new message received, said step of determining if all
5 members of said group has received the logged message includes:
6 deriving acknowledgments from state information included in
7 said gossip messages; and,
8 verifying receipt of the logged message for all members in
9 said pending acknowledgment list.